

UNCLASSIFIED

AFIT/EN/TR96-06

Air Force Institute of Technology

Probabilistic Temporal Networks

Eugene Santos Jr. Joel D. Young
esantos@afit.af.mil jdyoung@afit.af.mil

November 13, 1996

19970521 021

Approved for public release; distribution unlimited

19970521 021

Probabilistic Temporal Networks

Eugene Santos Jr.¹ and Joel D. Young

Department of Electrical and Computer Engineering

Air Force Institute of Technology

Wright-Patterson AFB, OH 45433-7765

November 13, 1996

Abstract

Complex real-world systems consist of collections of interacting processes/events. These processes change over time in response to both internal and external stimuli as well as to the passage of time itself. Many domains such as real-time systems diagnosis, story understanding, and financial forecasting require the capability to model complex systems under a unified framework to deal with both time and uncertainty. Current models for uncertainty and current models for time already provide rich languages to capture uncertainty and temporal information respectively. Unfortunately, these semantics have made it extremely difficult to unify time and uncertainty in a way which cleanly and adequately models the problem domains at hand. This is further compounded by the practical necessity of efficient knowledge engineering under such a unified framework. Existing approaches suffer from significant trade offs between strong semantics for uncertainty and strong semantics for time. In this paper, we define and explore a new model, the Probabilistic Temporal Network, for representing temporal and atemporal information while fully embracing probabilistic semantics. The model allows representation of time constrained causality, of when and if events occur, and of the periodic and recurrent nature of processes. A constraint satisfaction formulation is presented for belief revision as well as a polynomial solvable class.

Keywords: Knowledge Representation, Uncertainty, Temporal Reasoning, Probabilistic Representation

¹This research was supported in part by AFOSR Project #940006 and #9600989.

1 Introduction

In the evolution of expert systems, many techniques have been developed to represent human knowledge. One of the earliest is to represent knowledge as a logical system of if-then style rules (rule-based systems [4, 10]). A more recent approach is to represent knowledge (including uncertainty) of a situation, or "domain," as a network of states and probabilities (Bayesian Networks [20]).

Many domains, whether they are rule-based, probabilistic, or other, require a representation of time and of the temporal relationships between events. Most systems rely on a mechanism in which a date is associated with each piece of knowledge. Relationships are then determined simply by the date ordering. In more complicated domains, such as emergency room diagnosis, the date mechanism is not sufficient; one must be able to represent situations with relative knowledge like "precedes" or "during."

Real-world domains requiring a unified model of time and uncertainty include dealing with real-time system diagnosis, story understanding, planning and scheduling as well as financial forecasting. For example, consider the following scenario found in computer security analysis:

The computer operations center has a secure vault with a time-coded lock. This time-lock allows the vault to be opened from 0900 hours to 0905 hours and from 2100 to 2105. The center has critical operations from 0855 to 1805. Access to the vault is needed during the day and during critical operation making the vault likely to be open at those times. However, if the vault is closed, it can not be reopened until the time-lock allows.

This provides a detailed description of the causal and temporal relationships necessary to properly model the secure vault. As part of the computer security analysis, we must be able to translate this description and capture the knowledge in a form which we can correctly process and reason over.

Once the knowledge representation is captured, inferences can be made. Inferences can be of several types including prediction and explanation. Prediction is concerned with extending forward from the known past and present to the unknown future (statistical syllogism [13]). Explanation involves the determination of causality by extending from known data back to hypotheses (abduction) [13].

Complex systems consist of collections of interacting processes. These processes change over time in response to both internal and external stimuli as well as to the passage of time itself. There is great variety in the behavior of processes. Some processes are simple events such as opening a door or flipping a switch. Others are complex. One example being a communication channel, in which errors may occur due to lightning strikes and are more likely to occur given previous errors. Processes can also be recurrent or periodic, such as the passing of day into night or shifts in a work schedule.

The problem is to develop a model capable of representing complex systems changing over time. Given evidence about the past and present state of a system, one must be able to predict the system's future state. Also, given a future state, one must be able to determine the most probable causes. As knowledge about

such systems is bound to be incomplete and as the systems themselves may not be deterministic, the model must be able to represent uncertainty. This uncertainty permeates all areas, the duration of events, the strength of causal influence, the precise temporal relationship between events, etc.

Bayesian networks [20] provide a robust, probabilistic method of reasoning with uncertainty. Bayesian networks, however, do not provide a direct mechanism for representing temporal dependencies. For example, it is difficult to represent a situation such as the variability of an employee's arrival at work and the causal relationships between the time of arrival and later events.

Prior temporal modeling techniques have made trade-offs in expressiveness between semantics for time and semantics for uncertainty. Significant research has been done exploring time nets (also called time-slice Bayesian networks) [15, 14, 11]. These approaches build on the strong probabilistic semantics of Bayesian networks for expressing uncertainty. The discrete time net approach developed by Kanazawa models time as a series of points [14]. Events are considered to occur at an instant of time while facts are considered to occur over a series of time points. Both events and facts are represented by random variables. If dependencies only connect between random variables at the same or consecutive time points; then the net is said to be a Markov time net. In other words, the Markov property holds for a model when the future is conditionally independent of the past, given the present [15].

Hanks et al [11] is especially interesting for our work due to the emphasis on both endogenous and exogenous change [11]. Endogenous change is triggered by internal action, such as the progression of disease, and exogenous change is triggered by external change such as the administration of drugs. In our model, individual processes within a system must be able to respond to both endogenous (internal) and exogenous (external) stimuli.

The time-sliced approaches mentioned above are based on point models of time and, as such, require that events occur instantaneously. Often it is more natural to consider events as taking place over intervals of time. Also, the relationships between events that occur over intervals can be quite difficult to represent with only the three point relations (precedes, follows, equals).

Santos' Temporal Abduction Problem (TAP) [24] uses an interval representation of time. In the TAP, each event has an associated interval during which the event occurs. Relationships between events are expressed as directed edges from cause to effect within a weighted and/or directed acyclic graph structure. Edges are qualified with the possible interval relations. This allows great flexibility in expressing the relationship between events. For example, if event A must occur either before or after event B then the relation is written $\{<, >\}$. The TAP is an extension of cost based abduction [6] using a numeric cost to indicate the uncertainty of an event's occurrence. These costs are generally determined in an ad hoc manner by the domain expert. The TAP trades strong semantics of uncertainty for a powerful and flexible temporal representation.

This paper presents a new model, the Probabilistic Temporal Network (PTN), for representing temporal and atemporal information while remaining fully probabilistic. The model allows representation of time constrained causality, of when

and if events occur, and of the periodic and recurrent nature of processes. Bayesian networks lie at the foundation of the system and provide the probabilistic basis. Allen's interval system [1] and his thirteen relations provide the temporal basis.

PTNs focus on directly modeling processes and the interaction between them. The state of a process is represented by a value at a given time interval. A process can be defined over any number of such intervals. Random variables from traditional probability theory are used to model a process' value over each time interval.

We first briefly discuss temporal reasoning and Bayesian networks. From this foundation, the theoretical structure of our model is developed and its probabilistic nature proven. A linear constraint system for performing belief revision is developed as well as a polynomial solvable subclass. Along the way, several examples are developed including the secure vault scenario introduced above.

2 Temporal Reasoning

Temporal reasoning has been defined as the ability to reason about the relationships in time between events [10]. It is necessary to reason about time in many domains including planning, simulation, natural language understanding, and diagnosis. Temporal reasoning has been considered in philosophy and logic since Thales and Zeno [17]; however, it is only in the last two decades that temporal reasoning has been explicitly considered in artificial intelligence. McDermott and Allen, with their work in the early eighties [1, 2, 3, 18], brought temporal reasoning into the AI mainstream. Other models for temporal reasoning include algebras [30], semi-intervals [9], temporal constraint networks [8], and weak representations of interval algebras [16].

McDermott provides one of the earliest temporal representations [18]. In his approach, time is divided into a series of states with each state having an associated date, i.e. point in time. Facts are expressed as being true during particular states.

Allen introduced interval temporal reasoning to the AI community [1, 3]. Allen's interval algebra is governed by 13 relations on the intervals. Each event has an associated interval, denoted $[a, b]$, where a is the starting time point and b is the termination point. Temporal relationships between events are expressed as relations between their intervals. The relations between intervals, denoted \mathcal{A} , are $\{=, <, >, m, mi, d, di, s, si, f, fi, o, oi\}$ [1] (see Table 1). For example, event $A = [a, b]$ preceding event $B = [c, d]$ is denoted $A < B$ indicating that $a < b < c < d$. These relations are mutually exclusive and exhaustive. Note that, while there are thirteen relations between intervals, only three relations exist between points: precedes, equals, and follows.

Of special importance is Allen's use of disjunctive sets to express uncertainty in the exact relationship between intervals. For example, "interval A precedes or meets interval B " is written as $A\{<, m\}B$. Some commonly used disjunctions

Symbol:	Name:	Relation:
=	=	equals
<	>	precedes
m	mi	meets
d	di	during
s	si	starts
f	fi	finishes
o	oi	overlaps

Table 1: The thirteen possible interval-interval relations.

are *disjoint*, written $\{<, >, m, mi\}$, and *contains*, written $\{di, si, fi\}$ [1]. This can be represented in a graphical form where nodes represent events and the arcs are labeled with a disjunction of relations. The goal is to determine whether there exists an interval assignment to all the events that satisfy the disjunctive relations. If such a solution exists, then the given knowledge base is consistent.

While there is debate, in both philosophy and artificial intelligence, as to which representation, points or intervals, is most appropriate; the expressive power of the two methods is generally considered equivalent [14, 1] as intervals can be represented with beginning and end points in a point based approach. Allen points out, however, some paradoxes that can occur when points are allowed as the fundamental unit of time [1]. The problems arise from the durationless nature of points. Durationless intervals are not allowed, i.e. for any interval $[t_1, t_2]$, $t_2 > t_1$. If $t_1 = t_2$ is allowed then the thirteen interval-interval relations are not mutually exclusive. For example $[t_1, t_2]$ *starts* $[t_2, t_3]$ is indistinguishable from $[t_1, t_2]$ *meets* $[t_2, t_3]$ when $t_1 = t_2$. Mathematically, point relations should be expressed as $t_1 \mathcal{R} [t_2, t_3]$ and as such, there is a different set of point-interval relations which would add unnecessary overhead if used in our model. Our model strictly adheres to the philosophy that intervals are primitive and have non-zero duration.

Definition 1. A temporal interval is a closed interval $[a, b]$ on the reals (rationals if countability is an issue) such that $a < b$.

Axiom 1. The temporal interval is the primitive temporal individual.

Since all intervals must have non-zero duration, how can point intervals be expressed? The standard approach is to use $[t_0, t_0 + \epsilon]$ where ϵ is arbitrarily close, but not equal, to zero. Note that ϵ can be either added to the end or subtracted from the beginning or both. This approach is adopted in the PTN. To facilitate specifying the relationships between intervals, ϵ is deemed constant across an entire model. Thus $[t_0, t_0 + \epsilon] \{m\} [t_0, t_1]$ does not hold while $[t_0, t_0 + \epsilon] \{m\} [t_0 + \epsilon, t_1]$ does.

Aside from the temporal domain, neither Allen's nor McDermott's method

can explicitly model uncertainty. Uncertainty arises from many sources including missing or unavailable data as well as over generalization of rules [10]. For example if we have the rule "Birds Fly" and "Ostriches are birds" we conclude that "Ostriches fly." To prevent such a conclusion, additional rules must be added such as "Some birds fly" or "Ostriches don't fly." These additional rules can add significant complexity to a knowledge base.

3 Bayesian Networks (BNs)

Approaches to dealing with uncertainty include fuzzy logic [32], cost based techniques [6], certainty factors [27, 28], Dempster-Shafer theory [25], and probabilistic methods [20]. These approaches can be used both extensionally and intensionally. Extensional systems, such as rule-based systems, attach some sort of truth value to each rule or formula. The truth-value for some formula is calculated functionally from the truth-value of sub-formulae. Intensional systems, such as model-based systems, attach uncertainty to the possible states of the system itself [20]. Extensional systems are generally computationally efficient but their uncertainty measures are semantically weak. Intensional systems, on the other hand, are generally computationally expensive and semantically strong [20]. By carefully restricting which parts of an intensional system are relevant to the other parts, the computational limitations can, to some degree, be overcome.

In probabilistic reasoning, random variables (RVs) are used to represent events and/or objects in the world. By assigning various values to these RVs, we can model the current state of the world and weight the states according to the joint probabilities.

Bayesian networks are probabilistic intensional systems in which independence assumptions are used to restrict relevance. A Bayesian network is a directed acyclic graph (DAG) of random variable (RV) relationships. Directed arcs between RVs represent conditional dependencies. When all the parents of a given RV are instantiated, that RV is said to be conditionally independent of the remaining, non-descendent RVs given knowledge of its parents. For a more formal description of the independence semantics in Bayesian networks, see *d-separation* and *I-maps* in Charniak [5] and Pearl [20]. Figure 1 presents a simple example of a Bayesian network.

In general, we are searching for the world state with highest likelihood. This is called *belief revision* [20]. Belief revision is best used for modeling explanatory/diagnostic tasks. Basically, some evidence or observation is given to us, and our task is to come up with a set of hypotheses that together constitute the most satisfactory explanation/interpretation of the evidence at hand. Belief revision is a form of *abductive reasoning* [12, 21, 6]. More formally, if W is the set of all RVs in our given Bayesian network and e is our given evidence², any complete instantiation to all the RVs in W that is consistent with e is called an

²That is, e represents a set of instantiations made on a subset of W .

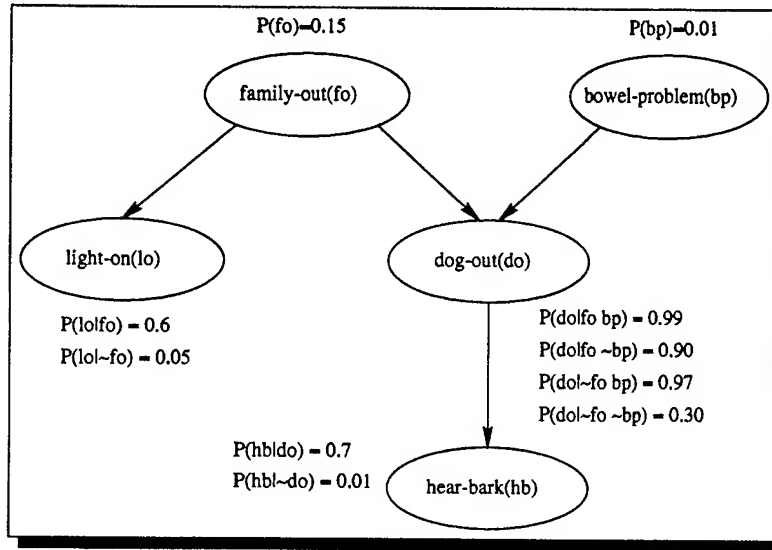


Figure 1: "Suppose when I go home at night, I want to know if my family is home before I try the doors. Now often when my wife leaves the house, she turns on an outdoor light. However, she sometimes turns on this light if she is expecting a guest. Also, we have a dog. When nobody is home, the dog is put in the back yard. The same is true if the dog has bowel troubles. Finally, if the dog is in the backyard, I will probably hear her barking." [5]

explanation or *interpretation* of e . The problem, then, is to find an explanation w^* such that

$$P(w^*|e) = \max_w P(w|e). \quad (1)$$

w^* is known as the *most-probable explanation*. The joint probability of any explanation w ,

$$w = (X_1 = x_1) \wedge (X_2 = x_2) \wedge \dots \wedge (X_m = x_m) \quad (2)$$

(where $X_1 \dots X_i \dots X_m$ is an arbitrary ordering of random variables in W , and x_i is some assignment to random variable X_i) is found using the chain rule [20]:

$$P(w) = P(x_m|x_{m-1}, \dots, x_1) \cdot P(x_{m-1}|x_{m-2}, \dots, x_1) \cdots P(x_2|x_1) \cdot P(x_1) \quad (3)$$

Bayesian networks take the chain rule one step further by making the important observation that certain RV pairs may become uncorrelated once information concerning other RV(s) is known. More precisely, we may have the following independence condition:

$$P(A|X_1, \dots, X_n, U) = P(A|X_1, \dots, X_n) \quad (4)$$

for some collection of RVs U . Intuitively, we can interpret this as saying that given knowledge of X_1, \dots, X_n knowledge of U is irrelevant to the state of A .

Combined with the chain rule, these conditional independencies allow us to replace the terms in the chain rule with smaller conditionals. Thus, instead of explicitly keeping the joint probabilities, all we need are smaller conditional probability tables, from which the joint probabilities can then be calculated.

For example, an application of the chain rule for computing the probability of an explanation for the Bayesian network in Figure 1 is

$$P(hb, do, lo, fo, bp) = \frac{P(hb|do, lo, fo, bp) \cdot P(do|lo, fo, bp)}{P(lo|fo, bp) \cdot P(fo|bp) \cdot P(bp)} \quad (5)$$

Using the dependencies we can simplify this to

$$P(hb, do, lo, fo, bp) = P(hb|do) \cdot P(do|fo, bp) \cdot P(lo|fo) \cdot P(fo) \cdot P(bp) \quad (6)$$

By choosing an ordering of the random variables consistent with the structure of the graph, such as that used in Equation 6 above, the savings from independencies is maximal and computation from the dependency tables in the Bayesian network is straightforward.

Bayesian networks [20] are a natural method for representing uncertainty. Bayesian networks, however, do not provide a direct mechanism for representing temporal dependencies. For example, it is difficult to represent a situation such as the variability of an employee's arrival at work and the causal relationships between the time of arrival and later events.

4 Combining Time and Probability

As previously discussed, the time-sliced approaches provide strong probabilistic semantics for representing uncertainty, however they are constrained in their temporal expressiveness. The TAP, on the other hand, has strong interval-based temporal semantics, but lacks strong probabilistic semantics.

What is needed, then, is a combined approach integrating strong probabilistic and temporal semantics. While much research has been done on point-based probabilistic temporal network models, little or no research has been identified using interval methods, specifically Allen's interval relations, for intensional probabilistic reasoning. As mentioned earlier, the interval representation of time is important for the expressive set of relations available. The closest research is the temporal abduction problem discussed above which does not have strict probabilistic semantics. Recent work by Young and Santos [31]³ does present a starting point, defining the network structure for a new model. The nodes of the network are temporal aggregates and the edges are the causal/temporal relationships between aggregates. Each aggregate represents a process changing over time. A temporal aggregate contains every interval of interest for the process. Each interval has an associated random variable giving the state of

³In which Probabilistic Temporal Networks (PTNs) are termed Temporal Bayesian Networks (TBNs) and Temporal Aggregates (TAs) are termed Temporal Random Variables (TRVs)

the process over that interval. Figure 2 depicts a temporal aggregate modeling when a vault is open. The 'Vault-Open' TA is dependent on itself (*VO*) and two other processes (*TU* and *CO*). This example is expanded into a full network below.




Figure 2: A simple temporal aggregate, 'Vault-Open', defined over four intervals. The conditional probability tables show 'Vault-Open' to be dependent on itself through some temporal causal relationship.

As is the case in the real world, the apparent state of a process is dependent on the temporal perspective of observation. An observation made in the middle of the night as to whether or not someone is at work may return different results than if the observation is made during the day. A switch can be turned on only if, at some previous time, the switch was turned off; the light can be on only when the switch is on.

To model the difference perspective makes in the apparent state of a process, edges in the network consist of a disjunctive set of interval relations and a schema to map the random variables of the intervals to a single value. This allows the exact definition of those intervals during which the state of one process affects another.

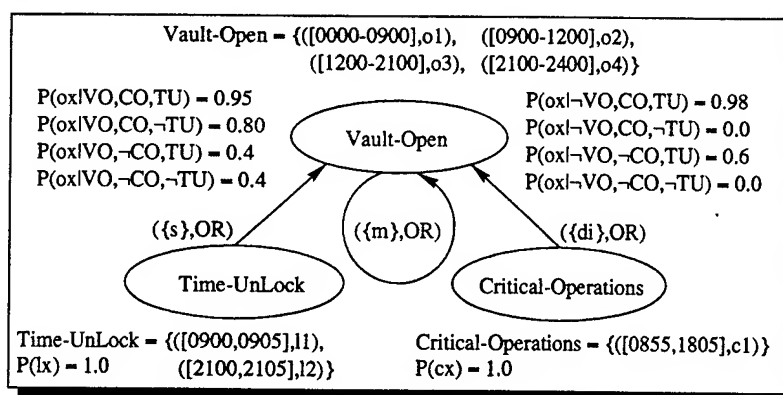


Figure 3: A probabilistic temporal network modeling a secure vault. This extends the 'Vault-Open' temporal aggregate in Figure 3.

Figure 3 shows a probabilistic temporal network modeling our earlier secure vault scenario detailing the various components and their interactions.

4.1 Temporal Aggregates

A process, such as ‘Vault-Open’ in Figure 3, is represented in the PTN by a temporal aggregate. Intuitively, a temporal aggregate consists of the set of states, e.g. $\{\text{true}, \text{false}\}$, $\{1, 2, 3\}$, or $\{\text{false}\} \cup \{\text{Red}, \text{Blue}\}$, that the process can take on, and a set of temporal intervals each having an associated random variable. Each such RV has a conditional probability table defined over the states of the process.

Definition 2. A temporal aggregate (TA) is an ordered pair (T, Σ) in which Σ is a set of states and T (pronounced Tau) is a set of ordered pairs (i, r) where i is a temporal interval and r is a random variable defined over Σ . For all pairs (i_1, r_1) and (i_2, r_2) in T , $r_1 = r_2$ iff $i_1 = i_2$. The dependencies for each random variable in the TA are defined only by temporal causal relationships between TAs.

In our prior work [31], temporal aggregates (then termed TRVs) were allowed to have internal dependencies to model endogenous change. This was found to be a source of temporal inconsistency and better represented through self loops as demonstrated in Figure 3. Endogenous change is explicitly modeled in the PTN with cyclic temporal causal relationships. This can be seen in the ‘Vault-Open’ process in Figure 3 in which the vault is more likely to stay open, given that it is open. Also note that this definition allows T to contain a potentially infinite number of interval-RV pairs. This paper assumes that temporal aggregates are finite, both in T and in Σ .

‘Vault-Open’ is formally written, according to Definition 2, as $VO = \{T, \Sigma\}$ where

$$T_{VO} = \{([0000, 0900], o_1), ([0900, 1200], o_2), ([1200, 2100], o_3), ([2100, 2400], o_4)\}$$

and

$$\Sigma_{VO} = \{\text{true}, \text{false}\}$$

with the conditional probability table being

$P(o_x VO, CO, TU)$	= 0.95	$P(o_x \neg VO, CO, TU)$	= 0.80
$P(o_x VO, CO, \neg TU)$	= 0.80	$P(o_x \neg VO, CO, \neg TU)$	= 0.0
$P(o_x VO, \neg CO, TU)$	= 0.4	$P(o_x \neg VO, \neg CO, TU)$	= 0.6
$P(o_x VO, \neg CO, \neg TU)$	= 0.4	$P(o_x \neg VO, \neg CO, \neg TU)$	= 0.0

for all RVs o_x where $o_x \in \{o_1, o_2, o_3, o_4\}$. The \neg symbol (as in $\neg TU$ above) indicates that the RV is assigned false, a non-negated RV (as in TU) indicates that the RV is assigned true.

Since $\Sigma = \{\text{true}, \text{false}\}$, $P(\neg o_x | VO, CO, TU) = 1 - P(o_x | VO, CO, TU)$. This holds for the other probabilities as well. In general, we will not explicitly show the probabilities when the true case is zero; e.g. $P(o_x | \neg VO, \neg CO, \neg TU) = 0.0$ would not be shown. Symbols used for temporal aggregates are uppercase letters from the end of the alphabet, e.g. X or Y , or uppercase abbreviations from the text name of the process being modeled, e.g. process ‘Vault-Open’ has a temporal aggregate denoted VO . Random variables within temporal aggregates

are denoted with lowercase letters, e.g. a , b , and c or y_1 and y_2 . Since the possible states of the aggregate are often evident from the conditional probability tables, Σ is often not explicitly shown. To differentiate between components of different temporal aggregates, the symbol of the component can contain the subscripted symbol of the associated TA, e.g. Σ_{VO} or σ_{1VO} .

An assignment to a temporal aggregate consists of an assignment to each interval-RV pair.

Definition 3. *A is an aggregate assignment (AA) iff A is a set of ordered pairs (τ, σ) where $\tau \in T$ and $\sigma \in \Sigma$ such that $\forall \tau \in T$, there exists a unique $\sigma \in \Sigma$ such that $(\tau, \sigma) \in A$. In other words, an aggregate assignment is a function from T into Σ .*

For example,

$$A_{VO} = \left\{ \begin{array}{l} ([0000, 0900], \text{false}), ([0900, 1200], \text{true}), \\ ([1200, 2100], \text{true}), ([2100, 2400], \text{false}) \end{array} \right\}$$

is an AA for the temporal aggregate VO from Figure 2. A_{VO} might be read "The vault was closed from 0000 hours to 0900 hours, open from 0900 hours to 2100 hours, and closed from 2100 hours to 2400 hours." The use of past tense here is arbitrary, *is closed* or *will be closed* would be equally appropriate. Aggregate assignments are denoted by uppercase letters from the beginning of the alphabet, e.g. A or B , subscripted if necessary by the symbol for the associated temporal aggregate.

Sometimes the entire state of a TA is not known. For example, we may only know that the vault was closed from 0000 to 0900. To express this we use a partial aggregate assignment which is simply a subset of an aggregate assignment.

Definition 4. *P is a partial aggregate assignment (PAA) for some temporal aggregate, X, iff there exists an A such that $P \subseteq A$ where A is an aggregate assignment for X. In other words, a partial aggregate assignment is a partial function from T into Σ .*

Our example, where the vault is only known to be closed over one interval is thus written:

$$P_{VO} = \{([0000, 0900], \text{false})\}$$

Note that P_{VO} is a subset of aggregate assignment A_{VO} above. PAAs are usually denoted by capital letters from the middle of the alphabet; however, since, by definition an aggregate assignment is also a PAA, some uppercase letters from the beginning of the alphabet may sometimes be used for PAAs.

4.2 Temporal Causal Relationships

How are the aggregates interconnected? The example network in Figure 4 shows a directed edge from 'Line-Open' to itself labeled $(\{m, o\}, \mathbf{OR})$. The edge combined with the conditional probability tables enforce a mutual exclusion constraint on 'Line-Open'. The communication line can be opened only if the line




Figure 4: A simple, one process probabilistic temporal network enforcing a mutual exclusion relationship. A communication line can only be opened given that it has not previously been opened.




Figure 5: The probabilistic temporal network from Figure 4 broken out to explicitly show the intervals (small circles) and the temporal relationships between intervals (dotted lines).

was not previously opened. Edges in the probabilistic temporal network are temporal causal relationships or TCRs.

While portrayed graphically as a labeled edge between temporal aggregates, the TCR is actually shorthand for a set of induced random variables that enforce the temporal constraints. These random variables combine the intervals selected by a disjunctive set of interval relations, e.g. $\{o, m\}$, using the probability distribution specified by a schema, e.g. **OR**, **XOR**, **PASSTHROUGH**. Figure 5 shows our example network from Figure 4 with the intervals and temporal relations explicitly shown. For example, the dotted line from interval lo_1 to interval lo_2 shows that lo_1 overlaps lo_2 . Figure 6 shows the network with the TCR replaced by the appropriate induced RVs.

What are the semantics behind the temporal causal relationship? The probability of some TA Y taking on some particular state over each interval is




Figure 6: The network in Figure 5 with the temporal causal relation replaced with the TCR induced random variables.

dependent on TA X taking on some state on interval(s) fitting the temporal relation, e.g. “no interval in Y can have state true unless that interval is after some interval in X having state true.” This is written $X(\{<\}, \mathbf{OR})Y$ with every $(i, r) \in T(Y)$ having conditional probabilities of the form $P(r | \dots, \neg X) = 0.0$. Schemas in general and the **OR** schema in particular are further discussed below.

Definition 5. A temporal causal relationship (TCR) describes a relationship between two temporal aggregates $X = (T_X, \Sigma_X)$ and $Y = (T_Y, \Sigma_Y)$ where X is considered the “cause” and Y the “effect.” Textually, the TCR is written $X(\mathcal{R}, \mathcal{M})Y$ where \mathcal{R} is a nonempty set of interval relations and \mathcal{M} is a schema for describing random variables. Graphically, the TCR is presented as a directed edge from the node for X to the node for Y , labeled with $(\mathcal{R}, \mathcal{M})$. Formally, the relationship is written as the four-tuple $(\mathcal{R}, \mathcal{M}, X, Y)$.

The TCR induces, for each interval-RV pair, (i_Y, r_Y) in T_Y , a random variable \mathcal{M}_r , defined over Σ_X , such that

1. r_Y is directly dependent on \mathcal{M}_r .
2. for each $(i_X, r_X) \in T_X$ where $i_X \mathcal{R} i_Y$, \mathcal{M}_r is directly dependent on r_X .
3. for each random variable x such that \mathcal{M}_r is directly dependent x , there exists an i_X such that $(i_X, x) \in T_X$.
4. the conditional probability table for \mathcal{M}_r is defined by the schema \mathcal{M} .

Temporal causal relationships are rarely given explicit names. Notationally, the random variables in the interval-RV pairs in the effect TA are usually written, in the conditional probability tables, as being dependent simply on the cause TA. This can be seen in the tables for the ‘Vault-Open’ temporal aggregate in Figure 3. In cases where there is more than one TCR between two TAs, some appropriate name or symbol can be associated with the TCR and the dependencies in the effect TA can be written as the name of the cause TA subscripted with the name of the TCR.

The random variable schema algorithmically defines the conditional probability tables for the random variables induced by the temporal causal relationship.

Definition 6. A random variable schema \mathcal{M} takes as parameters a set of states Σ , a set of interval-RV pairs T with RVs defined over Σ , a single interval-RV pair (i, r) , and an algorithm Δ which together define the conditional probability table for a random variable \mathcal{M}_r with states Σ such that for each $(i_T, r_T) \in T$, \mathcal{M}_r is directly dependent on r_T . \mathcal{M}_r is directly dependent on nothing else. The conditional probability table for \mathcal{M}_r is constructed with an algorithm, Δ . Δ can be either declarative or procedural.

For many models, these schemas are extremely simple, e.g.

$$\text{OR} : \left(\begin{array}{c} T, \\ \Sigma = \{\text{true}, \text{false}\}, \\ (i, r), \\ \Delta_{\text{OR}} \end{array} \right) \rightarrow \text{OR}_r$$

where Δ_{OR} is defined as

Algorithm 1: (Δ_{OR})

1. Let $(i_{T_1}, r_{T_1}) \dots (i_{T_n}, r_{T_n})$ be an arbitrary ordering of the elements of T
2. Create random variable OR_r such that \forall assignments A to $\{r_{T_1}, \dots, r_{T_n}\}$
 - (a) If there exists an $r \in A$ such that $r = \text{true}$

$$\begin{array}{l} \text{Let } P(\text{OR}_r = \text{true} | A) = 1 \\ \quad P(\text{OR}_r = \text{false} | A) = 0 \end{array}$$
 - (b) else
$$\begin{array}{l} \text{Let } P(\text{OR}_r = \text{true} | A) = 0 \\ \quad P(\text{OR}_r = \text{false} | A) = 1 \end{array}$$

Exclusive-or, **XOR**, can be defined by changing “there exists an $r \in A$ ” in step 2a above to “there exists a unique $r \in A$.” The other logical operations are also easily defined.

The schema **PASSTHROUGH**, defined:

$$\text{PASSTHROUGH} : \left(\begin{array}{c} T = (i_T, r_T), \\ \Sigma, \\ (i, r), \\ \Delta_{\text{PASSTHROUGH}} \end{array} \right) \rightarrow \text{PASSTHROUGH}_r$$

with $\Delta_{\text{PASSTHROUGH}}$ defined as

Algorithm 2: ($\Delta_{\text{PASSTHROUGH}}$)

1. Create random variable PASSTHROUGH_r such that for each $\sigma \in \Sigma$

$$\begin{array}{l} P(\text{PASSTHROUGH}_r = \sigma | r_T = \sigma) = 1 \\ P(\text{PASSTHROUGH}_r \neq \sigma | r_T = \sigma) = 0 \end{array}$$

produces a random variable for a causal relationship from a singleton TA (only one interval-RV pair in T). The temporal causal relationship

$$X(\mathcal{A}, \text{PASSTHROUGH})Y,$$

read “ X exerts direct causal influence on Y under all temporal relationships” is analogous to the causal relation in Bayesian networks. This type of relationship is useful when ‘temporalizing’ existing Bayesian networks.

4.3 Probabilistic Temporal Networks

A probabilistic temporal network is a directed graph in which the nodes are TAs and the edges are temporal causal relationships.

Definition 7. A probabilistic temporal network (PTN) is an ordered pair, (R, E) , where R is a set of temporal aggregates and E is set of temporal causal relationships such that, for each TCR in E from some temporal aggregate, X , to some temporal aggregate, Y , both X and Y are in R .

If each temporal aggregate in a probabilistic temporal network is assigned, then that PTN is said to be completely assigned. The set of all of the assignments and associated temporal aggregates forms a complete assignment.

Definition 8. The set \mathcal{C} containing (temporal aggregate, aggregate assignment) pairs is a complete assignment (CA) of some PTN (R, E) iff

1. $\forall (X, A) \in \mathcal{C}, X \in R$ and A is an aggregate assignment of X .
2. $\forall (X, A), (Y, B) \in \mathcal{C}, X = Y \Rightarrow A = B$.
3. $\forall X \in R \exists (Y, A) \in \mathcal{C}$ such that $X = Y$.

Complete assignments are denoted by uppercase script letters from the beginning of the alphabet, e.g. \mathcal{A}, \mathcal{B} , or \mathcal{C} .

When inferencing over a probabilistic temporal network, incomplete evidence as to the state of the network may be held. Such evidence is represented with a partial assignment. In the simplest form, any subset of a complete assignment is a partial assignment. A more complicated case arises when only a partial aggregate assignment is known for some temporal aggregate. Since a PAA is a subset (possibly improper) of an aggregate assignment, a partial assignment to a PTN consists of a subset of the variables of the PTN and associated partial aggregate assignments for the TAs. More formally:

Definition 9. The set \mathcal{P} containing (temporal aggregate, aggregate assignment) pairs is a partial assignment (PA) of some PTN (R, E) iff

1. $\forall (X, P) \in \mathcal{P}, X \in R$ and P is a partial aggregate assignment of X .
2. $\forall (X, P), (Y, Q) \in \mathcal{P}, X = Y \Rightarrow P = Q$.

PAs are usually denoted with uppercase script letters from the middle of the alphabet, e.g. \mathcal{P} or \mathcal{Q} . As a complete assignment is a subset of itself, by definition any complete assignment is also a partial assignment.

Notation. A partial assignment, \mathcal{P} , is said to be a subset of another partial assignment, \mathcal{Q} , (denoted $\mathcal{P} \subseteq \mathcal{Q}$) if every (X, P) in \mathcal{P} (except those having $P = \emptyset$) has a corresponding (Y, Q) in \mathcal{Q} such that $X = Y$ and $P \subseteq Q$. A complete assignment, say \mathcal{C} , is said to be compatible with a partial assignment, \mathcal{P} , if $\mathcal{P} \subseteq \mathcal{C}$, otherwise \mathcal{C} is said to be incompatible with \mathcal{P} . If \mathcal{C} is incompatible with \mathcal{P} then at least one temporal aggregate in \mathcal{C} has a different assignment than that in \mathcal{P} .

The goal of belief revision is to find the most probable state of the world given some evidence. This is the *most probable explanation*.

Definition 10. Let B be a PTN, let \mathcal{P} be partial assignment (evidence) of B , and let \mathcal{C} be some complete assignment (explanation) of B . \mathcal{C} is a most probable explanation (MPE) given \mathcal{P} iff for all \mathcal{A} where each \mathcal{A} is a complete assignment of B compatible with \mathcal{P} , $P(\mathcal{C}|\mathcal{P}) \geq P(\mathcal{A}|\mathcal{P})$.

Since $P(\mathcal{A}|\mathcal{P}) = P(\mathcal{A}, \mathcal{P})/P(\mathcal{P})$ and an incompatible complete assignment can not be a MPE (unless the evidence \mathcal{P} is itself contradictory in which case all CAs are MPEs), we only need to consider as candidates those complete assignments for which $\mathcal{P} \subseteq \mathcal{A}$. Thus since $\mathcal{P} \subseteq \mathcal{A}$, we derive $P(\mathcal{A}|\mathcal{P}) = P(\mathcal{A})/P(\mathcal{P})$. Furthermore, since $1/P(\mathcal{P})$ is a factor in the conditional probability of each explanation \mathcal{A} , to find the MPE, we need only compute the probability of each complete assignment, i.e. $P(\mathcal{A})$. $P(\mathcal{A})$ is calculated with the chain rule.

5 Cycles and Temporal Ordering

Now that the basic definitions and properties have been introduced, we will briefly explore the probabilistic temporal network in Figure 4 and consider a potential alternate representation. Figure 4 shows a network using a cyclic dependency to represent the internal dependencies in process 'Line-Open', i.e., a cyclic TCR has been used to explicitly model the endogenous temporal relationships. For 'Line-Open' to be true over some interval, 'Line-Open' must not be true over any earlier intervals.

Examining the intervals, "earlier" turns out to be either *meets* or *overlaps*. This is represented with a disjunctive set containing *meets* and *overlaps*: $\{m, o\}$. The conditional dependencies are represented using the **OR** schema. The TCR, $LO(\{m, o\}, \mathbf{OR})LO$, describes the random variable \mathbf{OR}_{lo_3} such that $P(\mathbf{OR}_{lo_3}|\neg lo_1, \neg lo_2) = 0$ and $P(\neg \mathbf{OR}_{lo_3}|\neg lo_1, \neg lo_2) = 1$. \mathbf{OR}_{lo_3} replaces LO in $P(lo_3|\neg LO) = 1$ to yield $P(lo_3|\neg \mathbf{OR}_{lo_3}) = 1$. By using cyclic TCRs to explicitly represent the temporal relationships within a process, the knowledge engineer can more clearly "see" the nature of the system being modeled.




Figure 7: The network in Figure 4 rewritten using a cyclic dependency such that the conditional probability table for each RV can be written with the same probability 1 instead of the dependent probabilities 1/3, 1/2, and 1 (not well-formed).

Figure 7 shows an attempt to simplify the conditional dependencies in process 'Line-Open'. The conditional probability tables for each random variable in process LO are identical. This is accomplished using the TCR $LO(\mathcal{A} - \{=$

$\}, \text{OR})LO^4$, which states that the random variable in each interval-RV pair is dependent on the random variables in all the other interval-RV pairs. While visually similar to the network in Figure 4, there is a serious problem with this network.




Figure 8: Process 'Line-Open' from Figure 7 drawn with the TCR $LO(\mathcal{A} - \{=\}, \text{OR})LO$ expanded. The loop shows a cycle in the dependencies.

Figure 8 shows process 'Line-Open' with the TCR expanded into the induced random variables. Notice that this violates the conditional independence assumptions discussed in the presentation of Bayesian networks. Random variable lo_2 is dependent on OR_{lo_2} which is dependent on lo_1 which is dependent on OR_{lo_1} which is dependent on lo_2 which is \dots lo_2 is separated from itself by random variables OR_{lo_2} , lo_1 , and OR_{lo_1} indicating that given knowledge of each of these variables that lo_2 is independent of itself which is clearly contradictory.

Figure 4 demonstrates an example in which a cycle in the PTN provided a useful representation of the internal dependencies within a process. Figure 7, on the other hand, shows a case in which the cycle, while intuitively satisfying, violates the requirements of conditional independence. This raises the question: "Under what circumstances are cycles appropriate in probabilistic temporal networks?"

Definition 11. An expanded probabilistic temporal network (EPTN) is the directed graph created by expanding all temporal causal relationships in some PTN.

Figure 9 shows the expanded probabilistic temporal network for the PTN from Figure 3. The OR node for o_1 is not shown as it has no parents and does not affect the probability distribution, i.e., $P(\text{OR}_{o_1} = \text{false}) = 1.0$. Note that a given EPTN is not necessarily a Bayesian network. Cycles can exist or extraneous arcs can be present, i.e., not a minimal I -map. Redundant induced RVs may also be present. Figure 10 presents an optimized network with an equivalent joint distribution as that of Figure 9. This optimization process is an avenue of further research.

⁴The set, $\mathcal{A} - \{=\}$, consists of all thirteen interval relations sans equals




Figure 9: Expanded Probabilistic Temporal Network for PTN in Figure 3. Labels on arcs indicate temporal relation; during inverse, starts, meets.

Definition 12. A probabilistic temporal network is said to be well-formed iff the corresponding expanded probabilistic temporal network contains no directed cycles, i.e., the EPTN of a well-formed PTN is a directed acyclic graph.

Figure 8, shown previously, gives an example EPTN with cycles. As discussed, this is problematic. A well-formed probabilistic temporal network does not contain any such directed cycles.

Since a temporal aggregate can contain an infinite number of interval-RV pairs, clearly the EPTN of an arbitrary PTN may be infinite. For the remainder of the paper we will assume that we are dealing only with finite networks.

Lemma 1 ([20]). For any DAG D there exists a probability distribution P such that D is a perfect map of P relative to d -separation, i.e., P embodies all the independencies portrayed in D , and no others.

This lemma, combined with Definition 12, leads directly to

Theorem 1. For each well-formed, finite PTN (R, E) there exists a probability distribution P such that P embodies all the independencies in (R, E) , and no others.

Theorem 1 indicates that if we have a well-formed, finite PTN, then we have an associated probability distribution. How can we guarantee that a given PTN is well-formed and finite? If there are a finite number of temporal aggregates in the PTN and each aggregate contains only a finite number of interval-RV pairs, then the PTN is finite. As mentioned earlier, this is assumed. Clearly if the PTN structure itself contains no cycles then there can be no cycles in the EPTN and our PTN is well-formed. The problem with this is that we lose significant expressive power. Networks such as that in Figure 3 would not be allowed.

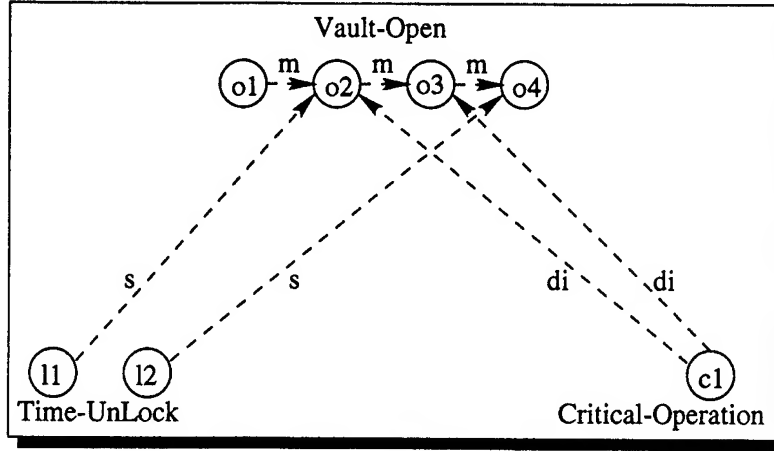


Figure 10: Optimized network for EPTN in Figure 9.

Cycles in the EPTN occur when an interval-RV pair becomes self-dependent. If only temporal relations which are strictly one directional are used, an interval-RV pair can not possibly be self-dependent. For example, if only $\{<\}$ is used in a PTN, no cycles are possible. The authors, in the development of the temporal abduction problem, defined the concept of *monotonicity* [23] as applied to temporal relations.

Definition 13. A set \mathcal{R} of temporal relations is said to be monotonic if and only if for all R in \mathcal{R} , $\overline{R}^c \cap (\overline{R}^c)^{-1} = \emptyset$ where $\overline{R} = \cup_{R \in \mathcal{R}} R$ and \overline{R}^c is the transitive closure of R and \overline{R}^{c-1} is the inverse of the transitive closure of R .

In the same work, we introduced the following monotonic set:

Proposition 1. The subset of relations $\mathcal{C} = \{<, o, s, fi, di, m\}$ from the original thirteen is a monotonic set.

Intuitively, a monotonic set, such as \mathcal{C} above, can be said to temporally ‘point in only one direction.’ This is compatible with Suppes’ probabilistic theory of causality [29] and Shoham’s criteria for causation [26] (both point based approaches) in which causation can only extend forward in time. For this reason, \mathcal{C} is said to be the *causal* set of temporal relations. The network in Figure 4 holds to \mathcal{C} .

Theorem 2. If, for probabilistic temporal network (R, E) , there exists a monotonic set, \mathcal{Q} , of temporal relations such that for each $(R, \mathcal{M}, X, Y) \in R$, $R \subseteq \mathcal{Q}$; then the PTN (R, E) is well-formed.

Proof. sketch Since the only temporal relations used in the PTN are drawn from \mathcal{Q} and \mathcal{Q} is *monotonic*, no interval-RV pair can ever relate to itself temporally (otherwise $\overline{\mathcal{Q}}^c \cap (\overline{\mathcal{Q}}^c)^{-1} \neq \emptyset$) and as there can be no cycles within the TAs themselves; there can be no cycles in the EPTN and thus the PTN is well-formed. \square

Definition 14. A causal probabilistic temporal network (CPTN) is a PTN for which Theorem 2 holds with $\mathcal{R} = \mathcal{C}$.

Figure 1 is a state transition diagram for a two-person conversation. The diagram consists of three states: 'Trigger', 'Person B Talking', and 'Person A Talking'. The transitions are as follows:

- From 'Trigger' to 'Person B Talking': $P(S_i) = 0.2$ ($0 < i < n/2$)
- From 'Trigger' to 'Person A Talking': $P(S_i) = 0.2$ ($0 < i < n/2$)
- From 'Person B Talking' to 'Person A Talking': $P(B_i|T, A) = 0.85$ ($0 < i < n$)
- From 'Person A Talking' to 'Person B Talking': $P(A_i|T, B) = 0.85$ ($0 < i < n$)
- From 'Person B Talking' to 'Trigger': $P(B_i|T, \sim A) = 0.20$ ($0 < i < n$)
- From 'Person A Talking' to 'Trigger': $P(A_i|T, \sim B) = 0.20$ ($0 < i < n$)

The diagram also includes the following definitions:

- $S = \{ \{ (2 \cdot Ti - e, 2 \cdot Ti), Si \} \mid 0 < i < \text{ceil}(n/2) \}$
- $B = \{ \{ (Ti, Ti+1), Bi \} \mid 0 < i < n \}$
- $A = \{ \{ (Ti, Ti+1), Ai \} \mid 0 < i < n \}$

The diagram is a Petri net with the following components and connections:

- Places:** b_0, b_1, b_2, b_3 (top row); a_0, a_1, a_2, a_3 (bottom row); t_0 (left); t_1 (right).
- Transitions:** $OR_1, PT_1, OR_2, PT_2, OR_3, PT_3$ (middle row).
- Connections (dashed lines with arrows):**
 - $t_0 \rightarrow b_0$ and $t_0 \rightarrow a_0$
 - $b_0 \rightarrow OR_1$
 - $a_0 \rightarrow PT_1$
 - $OR_1 \rightarrow b_1$
 - $PT_1 \rightarrow a_1$
 - $b_1 \rightarrow OR_2$
 - $a_1 \rightarrow PT_2$
 - $OR_2 \rightarrow b_2$
 - $PT_2 \rightarrow a_2$
 - $b_2 \rightarrow OR_3$
 - $a_2 \rightarrow PT_3$
 - $OR_3 \rightarrow b_3$
 - $PT_3 \rightarrow a_3$
 - $b_3 \rightarrow t_1$ and $a_3 \rightarrow t_1$

The equals relation, $=$, is not a member of \mathcal{C} , and can not be a member of any *monotonic* set of relations as $=$ is its own inverse. Equals is, however, useful for expressing simultaneity. Figure 11 shows an example in which two people

are chatting. Talker *A* tends to ‘talk over’ Talker *B*. To model this, the TCR from *B* to *A* includes equals as well as meets. Figure 12 shows the EPTN for Figure 11.

To insure that a CPTN extended to use equals is well-formed, each directed cycle must have at least one TCR in which equals is not used. This guarantees ‘time progression’ in each cycle. A probabilistic temporal network limited to $\mathcal{C} \cup \{=\}$ with this broken cycle property is said to be S-Causal (SCPTN) (‘S’ for simultaneity).

6 Constructing a Partial Order and Using the Chain Rule

Previously we discussed finding the *most probable explanation*. The MPE is the complete assignment with the greatest joint probability. As mentioned, this joint probability is calculated using the chain rule. To efficiently use the chain rule, a partial ordering (from effect to cause) of the random variables must exist. The ordering is drawn from the expanded PTN and can only be found when the PTN is well-formed and finite. The following algorithm finds a partial ordering for a well-formed and finite PTN:

Algorithm 3: (Partial Ordering)

1. First, find the EPTN of a well-formed and finite PTN.
2. From the EPTN, select all RVs with no children. Place these first in the ordering in arbitrary order.
3. Find all RVs among all those not yet ordered such that all children thereof are ordered. Place these next in the ordering, again in arbitrary order.
4. Repeat Step 3 until no unordered RVs remain.

For example, the PTN in Figure 4 expands to the EPTN in Figure 6. A partial ordering of the RVs is found in the following steps:

1. Order: () RVs: $\{lo_1, lo_2, lo_3, OR_{lo_2}, OR_{lo_3}\}$
2. Order: (lo_3) RVs: $\{lo_1, lo_2, OR_{lo_2}, OR_{lo_3}\}$
3. Order: (lo_3, OR_{lo_3}) RVs: $\{lo_1, lo_2, OR_{lo_2}\}$
4. Order: (lo_3, OR_{lo_3}, lo_2) RVs: $\{lo_1, OR_{lo_2}\}$
5. Order: $(lo_3, OR_{lo_3}, lo_2, OR_{lo_2})$ RVs: $\{lo_1\}$
6. Order: $(lo_3, OR_{lo_3}, lo_2, OR_{lo_2}, lo_1)$ RVs: $\{\}$

yielding $(lo_3, OR_{lo_3}, lo_2, OR_{lo_2}, lo_1)$ as a partial ordering.

Since a partial ordering exists for the network, the chain rule can be used to find the joint probability of each assignment. Table 2 shows the probability distribution defined by the example in Figure 4. Only non-zero probability assignments are shown (but one).

Each joint probability in Table 2 is calculated using the chain rule [20]. For example, the probability of the complete assignment

$$\left\{ \left(LO, \left\{ \begin{array}{l} ([0900, 0910], lo_1), \text{true}), \\ ([0905, 0915], lo_2), \text{false}), \\ ([0910, 0920], lo_3), \text{false}) \end{array} \right\} \right\} \quad (7)$$

JOINT PROBABILITY TABLE FOR FIGURE 4					
Line-Open					Assignment
[0910,0920]	OR _{lo₃}	[0905,0915]	OR _{lo₂}	[0900,0910]	Probability:
true 1	false 1	false 1/2	false 1	false 2/3	1/3 (1)
false 1	true 1	true 1/2	false 1	false 2/3	1/3 (2)
false 1	true 1	false 1	true 1	true 1/3	1/3 (3)
true 0	true 1	false 1	true 1	true 1/3	0 (4)
Total:					1

Table 2: The possible complete assignments to the network in Figure 4 with associated probabilities. One ‘impossible’ assignment is also shown.

is calculated from

$$\begin{array}{rclcl}
P(lo_3 = \text{false} \mid \text{OR}_{lo_3} = \text{true}) & & 1 & & \\
P(\text{OR}_{lo_3} = \text{true} \mid lo_1 = \text{true}, lo_2 = \text{false}) & & 1 & & \\
P(lo_2 = \text{false} \mid \text{OR}_{lo_2} = \text{true}) & = & 1 & = & \frac{1}{3} \quad (8) \\
P(\text{OR}_{lo_2} = \text{true} \mid lo_1 = \text{true}) & & 1 & & \\
P(lo_1 = \text{true}) & & 1/3 & &
\end{array}$$

7 Constraint Satisfaction

In the previous section, we showed how to calculate the probability of a complete assignment to a probabilistic temporal network. In this section we present a method for finding the most probable complete assignment, i.e., performing belief revision on probabilistic temporal networks. We use a constraint satisfaction approach with mixed Boolean linear programming. Constraint satisfaction has three main advantages; first, constraints can be formed to take advantage of the inherent structure of the PTN; second, very efficient algorithms developed by the operations research community are available; and finally, alternate explanations, e.g. second or third best, can be found using techniques presented in [22].

Definition 15. A constraint system is a 3-tuple (Γ, I, ψ) where Γ is a finite set of variables, I is a finite set of linear inequalities based on Γ , and ψ is a cost function from $\Gamma \times \{\text{true}, \text{false}\}$ to \mathbb{R} .

Our probabilistic temporal network model can be considered to have a layered structure. The layers consist of temporal aggregates and temporal causal relationships. For this reason, we present our system of constraints in two parts,

those for TCRs and those for TAs. For some well-formed PTN $P = (R, E)$, the following steps produce the constraints, variables, and costs for the temporal causal relationships in E and those for the temporal aggregates in R , i.e. the following steps produce $L(P) = (\Gamma, I, \psi)$.

1. For each TCR $(\mathcal{R}, \mathcal{M}, (T_X, \Sigma_X), (T_Y, \Sigma_Y))$ in E ,

(a) For each $(i_Y, r_Y) \in T_Y$ construct variables $\mathcal{M}_{\sigma_{X_1}}^{r_Y} \dots \mathcal{M}_{\sigma_{X_n}}^{r_Y}$ in Γ where $\sigma_{X_1} \dots \sigma_{X_n}$ are states in Σ_X . Set costs for each variable as

$$\psi(\mathcal{M}_{\sigma_{X_i}}^{r_Y}, \text{false}) = \psi(\mathcal{M}_{\sigma_{X_i}}^{r_Y}, \text{true}) = 0. \quad (9)$$

where $1 \leq i \leq n$ and add the following constraint to I :

$$\sum_{i=1}^n \mathcal{M}_{\sigma_{X_i}}^{r_Y} = 1. \quad (10)$$

(b) $\forall (i_Y, r_Y) \in T_Y$ and each $\sigma_X \in \Sigma_X$ let $(i_{X_1}, r_{X_1}) \dots (i_{X_j}, r_{X_j}) \in T_X$ be those pairs for which $i_{X_h} \mathcal{R} i_Y$ with $1 \leq h \leq j$, then

i. for each conditional probability of the form

$$P(\mathcal{M}_{r_Y} = \sigma_X | r_{X_1} = \sigma_{X_1} \dots r_{X_j} = \sigma_{X_j})$$

as induced by schema \mathcal{M} , construct a variable

$$q[\mathcal{M}_{r_Y} = \sigma_X | r_{X_1} = \sigma_{X_1} \dots r_{X_j} = \sigma_{X_j}] \quad (11)$$

(denoted q in following steps) in Γ such that

A.

$$\psi(q, \text{false}) = 0, \quad (12)$$

$$\psi(q, \text{true}) = -\log \left(P \left(\mathcal{M}_{r_Y} = \sigma_X \left| \begin{array}{c} r_{X_1} = \sigma_{X_1} \\ \dots \\ r_{X_j} = \sigma_{X_j} \end{array} \right. \right) \right) \quad (13)$$

B. with the following constraint in I :

$$q \geq \sum_{h=1}^j T_{\sigma_{X_h}}^{r_{X_h}} + \mathcal{M}_{\sigma_X}^{r_Y} - j \quad (14)$$

(c) Let $\Upsilon_{\mathcal{M}_{\sigma_X}^{r_Y}}$ be the set of all q constructed in step (1b) for variable $\mathcal{M}_{\sigma_X}^{r_Y}$. For each such variable, add the following constraint to I :

$$\mathcal{M}_{\sigma_X}^{r_Y} = \sum_{q \in \Upsilon_{\mathcal{M}_{\sigma_X}^{r_Y}}} q. \quad (15)$$

2. For each TA $X = (T_X, \Sigma_X)$ in R

- (a) For each $(i_X, r_X) \in T_X$ construct variables $T_{\sigma_{X_1}}^{r_X} \dots T_{\sigma_{X_n}}^{r_X}$ in Γ where $\sigma_{X_1} \dots \sigma_{X_n}$ are states in Σ_X . Set costs for each variable as

$$\psi(T_{\sigma_{X_i}}^{r_X}, \text{false}) = \psi(T_{\sigma_{X_i}}^{r_X}, \text{true}) = 0. \quad (16)$$

where $1 \leq i \leq n$ and add the following constraint to I :

$$\sum_{i=1}^n T_{\sigma_{X_i}}^{r_X} = 1. \quad (17)$$

- (b) For each $(i_X, r_X) \in T_X$ and each $\sigma_X \in \Sigma_X$ let $\mathcal{M}_1 \dots \mathcal{M}_j$ be those random variables induced by TCRs $(\mathcal{R}_h, \mathcal{M}_h, Y_h, Z_h)$ for which $1 \leq h \leq j$ and $Z_h = X$. Then
- i. for each conditional probability of the form

$$P(r_X = \sigma_X | \mathcal{M}_1 = \sigma_{Y_1} \dots \mathcal{M}_j = \sigma_{Y_j}),$$

construct a variable

$$q[r_X = \sigma_X | \mathcal{M}_1 = \sigma_{Y_1} \dots \mathcal{M}_j = \sigma_{Y_j}] \quad (18)$$

(denoted q in following steps) in Γ such that

- A.

$$\psi(q, \text{false}) = 0, \quad (19)$$

$$\psi(q, \text{true}) = -\log \left(P \left(r_X = \sigma_X \left| \begin{array}{c} \mathcal{M}_1 = \sigma_{Y_1} \\ \dots \\ \mathcal{M}_j = \sigma_{Y_j} \end{array} \right. \right) \right) \quad (20)$$

B. with the following constraint in I :

$$q \geq \sum_{h=1}^j \mathcal{M}_{\sigma_{Y_h}}^{r_X} + T_{\sigma_X}^{r_X} - j \quad (21)$$

- (c) Let $\Upsilon_{T_{\sigma_X}^{r_X}}$ be the set of all q constructed in step (1b) for variable $T_{\sigma_X}^{r_X}$. For each such variable, add the following constraint to I :

$$T_{\sigma_X}^{r_X} = \sum_{q \in \Upsilon_{T_{\sigma_X}^{r_X}}} q. \quad (22)$$

In this construction, constraints (10) and (17) ensures that each random variable, either induced or in a TA, can take on one and only one value. Constraints (14) and (15) guarantee that each of the probabilities for TCR induced variables is computed in concordance with the appropriate temporal relations and schema. Constraints (21) and (22) guarantee that the probability of a temporal assignment to a TA is computed with the appropriate set of conditional probabilities. Variables of the form $q[r_X = \sigma_X | \mathcal{M}_1 = \sigma_{Y_1} \dots \mathcal{M}_j = \sigma_{Y_j}]$ are

called *conditional variables* in that they explicitly represent the dependencies between RVs and are the mechanism for computing the probability of any complete assignment.

For example, consider again the simple probabilistic temporal network in Figure 4. Previously we demonstrated how to calculate the probability of an assignment to this network using the chain rule (see Table 2 and Equation 8). Now, if we take the complete assignment

$$\left\{ \left(LO, \left\{ \begin{array}{l} ([0900, 0910], lo_1), \text{true}, \\ ([0905, 0915], lo_2), \text{false}, \\ ([0910, 0920], lo_3), \text{false} \end{array} \right\} \right) \right\} \quad (23)$$

we expect our variable assignments to be

$$\begin{aligned} LO_{true}^{lo_1} &= q[lo_1 = \text{true} | \text{OR}_{lo_1} = \text{false}] \\ LO_{false}^{lo_2} &= q[lo_2 = \text{false} | \text{OR}_{lo_2} = \text{true}] \\ LO_{false}^{lo_3} &= q[lo_3 = \text{false} | \text{OR}_{lo_3} = \text{true}] \\ \text{OR}_{false}^{lo_1} &= q[\text{OR}_{lo_1} = \text{false}] \\ \text{OR}_{true}^{lo_2} &= q[\text{OR}_{lo_2} = \text{true} | lo_1 = \text{true}] \\ \text{OR}_{true}^{lo_3} &= q[\text{OR}_{lo_3} = \text{true} | lo_1 = \text{true}, lo_2 = \text{true}] \end{aligned} = 1$$

with all other variables being zero. Since the only variables which incur costs are the $q[\dots]$ variables, the cost of the assignment is $-\log(1/3) - \log(1) - \log(1) - \log(1) - \log(1) - \log(1) = -\log(1/3)$ and thus the probability of the assignment is $1/3$ as expected. As informally demonstrated in this example, the cost of a variable assignment is found by summing the product of each variable in Γ and its corresponding cost in ψ .

Definition 16. A variable assignment for a constraint system $L = (\Gamma, I, \psi)$ is a function s from Γ to \mathbb{R} . Furthermore,

1. If the range of s is $\{0, 1\}$, then s is a 0-1 assignment.
2. If s satisfies all of the constraints in I , then s is a solution for L .
3. If s is a solution for L and is also a 0-1 assignment, then s is a 0-1 solution for L .

Definition 17. Given a constraint system $L = (\Gamma, I, \psi)$, we construct a function Θ_L from variable assignments to \mathbb{R} as follows:

$$\Theta_L(s) = \sum_{\gamma \in \Gamma} s(\gamma) \psi(\gamma, \text{true}) + (1 - s(\gamma)) \psi(\gamma, \text{false}) \quad (24)$$

Θ_L is called the objective function of L .

Definition 18. An optimal 0-1 solution for a constraint system $L = (\Gamma, I, \psi)$ is a 0-1 solution which minimizes Θ_L .

By finding an optimal 0-1 solution for a constraint system, we find the most probable explanation for the corresponding PTN. Santos [22] presents a customized algorithm using the *cutting plane method* [19] for finding the optimal 0-1 solution. Since any Bayesian network can be represented as a PTN⁵, we know that, in general, belief revision over PTNs is NP-hard [7, 20].

⁵Treat each RV in the BN as a TA with a single interval-RV pair, using the $\{=\}$, **PASSTHROUGH** TCR, and make all intervals in the TAs equivalent.

8 Polynomial Time Belief Revision—The Generalized Temporal Polytree

In the previous section, we presented a method for performing belief revision on probabilistic temporal networks. In general, this problem is NP-hard. However, for singly-connected PTNs (polytrees), belief revision can be done in polynomial time. A polytree is a directed acyclic graph in which no more than one path exists between any two nodes. The lack of undirected cycles in the graph structure allows for efficient local decisions. In this section we present the generalized temporal polytree (GTP); a PTN model with a restricted graph and temporal structure. The EPTN for a GTP is guaranteed to be a polytree.

First we introduce a pair of restrictions on the probabilistic temporal network. These two restrictions force the expanded PTN to be a causal tree, i.e., all nodes (except root nodes) have one and only one incoming edge (cause) ⁶. A causal tree structure allows for very easy belief updating and revision. The first requirement is that the only interval-interval relation allowed is *meets*. *Meets* enforces a strictly monotonic progress in time and, unlike *precedes*, does not allow “temporally remote causation [29].” The second requirement is that all intervals across the network have different end-points. Together, these two requirements impose a causal tree structure on the expanded network. A probabilistic temporal network holding to these two requirements is termed a Generalized Causal Temporal Tree.

Definition 19. A generalized causal temporal tree (*GCTT*) is a probabilistic temporal network in which

1. $\mathcal{R} = \{m\}$ for each $(\mathcal{R}, \mathcal{M}, X, Y) \in E$, i.e., *meets* is the only temporal relation allowed.
2. All intervals in all temporal aggregates must have unique end-points.

Theorem 3. The expanded probabilistic temporal network of any generalized causal temporal tree is a causal tree.

Proof. By Contradiction. Let $P = (R, E)$ be some generalized causal temporal tree. Let N be the EPTN of P . Assuming that N is not a tree, we know by the definition that there exists a node, a , such that at least two different directed edges enter a from two different causal nodes (ignoring intervening induced RVs), say b and c . Each of these nodes (a, b, c) have associated intervals, say, $([a_s, a_e], [b_s, b_e], [c_s, c_e])$ respectively. Since, by the definition of generalized causal temporal tree, $[b_s, b_e] \text{ meets } [a_s, a_e]$ and $[c_s, c_e] \text{ meets } [a_s, a_e]$; $b_e = a_s$ and $c_e = a_s$ and thus $b_e = c_e$. However, again from the definition of generalized causal temporal tree, all end-points are unique and thus $b_e \neq c_e$. \square

Corollary 1. The EPTN of a GCTT in which constraint 2 in Definition 19 is changed to start-points instead of end-points, has an inverted tree structure.

By connecting together regions with varying end-points (*out-regions*) with regions of varying start-points (*in-regions*) a PTN with polytree structure is

⁶Note that by this definition, we actually allow a collection of such unconnected trees.

formed. A region, then, is a collection of TAs in which all interval end or start points are different. Regions join together at a set of TAs referred to as a *join-region*. All TAs in a join-region are members of both regions being joined. For example, if an in-region and an out-region are joined, then all end-points in the join-region must be different from all end-points in the out-region and all start-points in the join-region must differ from all start-points in the in-region.

Definition 20. A set of temporal aggregates, R , forms an out-region if for each

$$([s_1, e_1], r_1) \in \bigcup_{(T, \Sigma) \in R} T$$

there does not exist another

$$([s_2, e_2], r_2) \in \bigcup_{(T, \Sigma) \in R} T$$

such that $r_1 \neq r_2$ and $e_1 = e_2$, i.e., all intervals in all temporal aggregates have unique end-points.

Definition 21. A set of temporal aggregates, R , forms an in-region if for each

$$([s_1, e_1], r_1) \in \bigcup_{(T, \Sigma) \in R} T$$

there does not exist another

$$([s_2, e_2], r_2) \in \bigcup_{(T, \Sigma) \in R} T$$

such that $r_1 \neq r_2$ and $s_1 = s_2$, i.e., all intervals in all temporal aggregates have unique start-points.

Definition 22. A set of temporal aggregates, R , forms a join-region for two in- or out-regions, R_1 and R_2 if $R = R_1 \cap R_2$

To prevent undirected cycles (directed cycles are prevented by the *meets* restriction), out-regions are not permitted to join to out-regions. In-regions can join with both in-regions and out-regions. No temporal causal relationships can extend, however, from a join-region back into an in-region. This prevents undirected cycles by enforcing the constraint that all inverted trees in an in-region must end in the join-region (or not enter the join-region).

Definition 23. A generalized temporal polytree (GTP) is a probabilistic temporal network $P = (R, E)$ for which there exist sets I (in-regions), O (out-regions), and J (join-regions) such that

1. $\mathcal{R} = \{m\}$ for each $(\mathcal{R}, \mathcal{M}, X, Y) \in E$, i.e., *meets* is the only temporal relation allowed.
2. Each TA in the PTN is in some in- or out-region and vice versa.
3. Each join-region in J connects two in-regions or connects an in-region with an out-region. Out-regions can not join with other out-regions.
4. For each TCR, $(\mathcal{R}, \mathcal{M}, X, Y) \in E$, exactly one of the following must hold:

- (a) there exists one and only one $r \in I \cup O$ such that $X, Y \in r$, or
- (b) there exists a $j \in J$ such that $X, Y \in j$, or
- (c) there exists a $j \in J$ and an $o \in O$ such that $X \in j$ and $Y \in o$.

In no case can X be in a join-region and Y be in an in-region outside of the join.

Theorem 4. *The expanded probabilistic temporal network of any generalized temporal polytree is a polytree.*

Proof. By Contradiction. Let $P = (R, E)$ be some generalized temporal polytree. Let N be the EPTN of P . Assuming that N is not a polytree, we know by definition that there exists at least two nodes such that two unique undirected paths exist between them. These two paths form an undirected cycle. Based on Theorem 3 and Corollary 1, there can not exist more than one unique path between any two nodes within any give in- or out-region. Also, different regions can only connect together in join-regions. Thus at least two nodes on the undirected cycle must be in the join-region. Let these two nodes be a and b . Since all nodes in the join-region belong to both in- or out- regions and no cycles can exist within any single in- or out- regions, at least one node on the cycle, say c , must exist outside of the join-region. This leads to two cases: either c is in an in-region or c is in an out-region. Either way if c is in one region and a and b in the join-region, there must be a fourth node, d , in the other region from c , otherwise the cycle would lie entirely within one in- or out-region.




Figure 13: Possible shape of an undirected cycle in a generalized temporal polytree.

This gives us four nodes on our cycle, a , b , c , and d . We know that a and b are both in the join-region and we know that both c and d are outside of the join-region and each in different regions. This gives us a structure as in Figure 13. Since out-regions can not join to out-regions, either node d or node c must lie in an in-region. Let us assume that this is node d . Since a TCR can not extend from the join-region out into an in-region, a TCR must extend from the TA containing d into the join-region. This TCR must be such that the interval associated with d meets two nodes in the join-region, however since all nodes in the join-region are also in the same in-region as d , no two nodes in the join-region can have the same start point and thus d can not meet these two nodes and thus an undirected cycle can not exist. \square

Although not stated in the formal definition of the generalized temporal polytree, interval start-points in evidence TAs and end-points in leaf TAs do not need to be different from other start- or end-points as evidence nodes are not dependent on anything and nothing is dependent on leaf nodes.

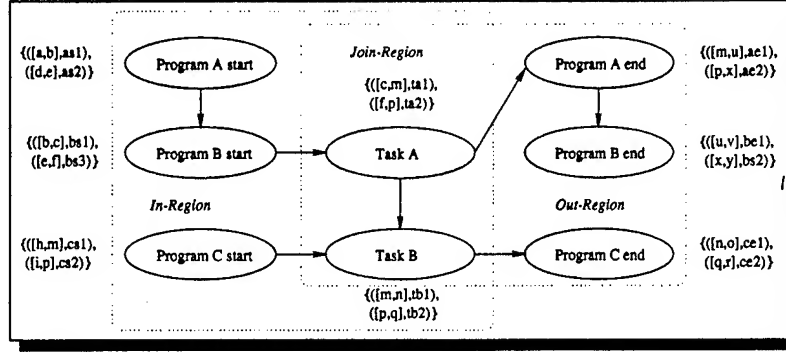


Figure 14: A Generalized Temporal Polytree depicting a program execution scenario.

Figure 14 shows a GTP modeling a program execution scenario. Program-A executes Program-B to complete Task-A. Program-C must complete Task-B. Task-B, however, requires that Task-A complete immediately prior. The start and task TAs form an in-region and the task and end TAs form an out-region. Task-A and Task-B together form a join-region. Figure 15 shows the expanded probabilistic temporal network for this GTP.

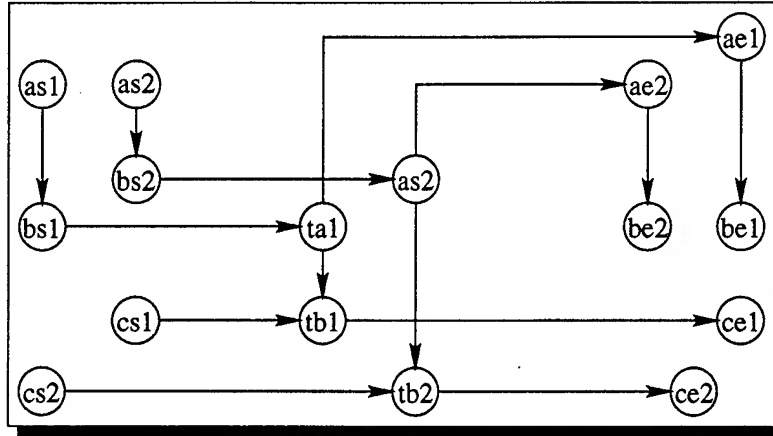


Figure 15: The EPTN for the GTP in Figure 14.

9 Related Work

In addition to the other work we have mentioned earlier, Aliferis and Cooper have also developed a preliminary temporally extended Bayesian network formulation termed the *Modifiable Temporal Bayesian Network-Single Granularity* (MTBN-SG). A MTBN-SG is primarily an extended time-sliced Bayesian network defined over a range of time points. Each *ordinary* node in a MTBN-SG is indexed over this entire range. Edges between nodes are represented by a *mechanism* variable which is a Boolean true/false random variable indicating whether the link is active, i.e. whether a dependency exists between the connected variables. Each such mechanism has an associated *lag* random variable (Delta TAs in the PTN) defined over the range of time points indicating the delay between the "cause" and the "effect." Atemporal or *abstract* random variable nodes are supported which are not instantiated for each time point. The resultant graph can have cycles to allow expressions of recurrence and feedback. As long as all cycles in the underlying joint distribution have zero probability, the graph is said to be *well-defined*.

Since the edges, both *mechanism* and *lag* components, are represented by random variables, the edges can be both dependent on and causal too other random variables in the network. This allows the knowledge engineer to express conditions where a relationship exists between variables only under certain circumstances. The problem with this approach is that joint distributions can be described which are not compatible with the Bayesian model. Maintaining consistency in the local probability tables across random variables then becomes a concern.

As indicated in the name, the MTBN-SG model only supports a single granularity for the size of the time step in any given network. Extending the model to support multiple granularities appears problematic, especially in the case when the granularities are not multiples, e.g. g_1 is every ten minutes and g_2 is every fifteen minutes. A, perhaps more difficult problem arises in the model if the start time for one granularity is not the same as that of another as the granularities may be forever out of phase. This is not an issue for our model. Individual processes or temporal aggregates can be modeled with arbitrary sets of intervals. There is no requirement that the intervals in one TA match those in other TAs as the temporal causal relationship describes the desired relationships.

Intervals can be modeled in the MTBN with *abstract* variables, *INT_START* and *INT_END*, representing the start and end points of the interval respectively. *INT_END* is dependent on *INT_START* such that the end time will never be before the start time. The duration of an interval can be acquired from a third variable, *INT_DUR* dependent on both *INT_START* and *INT_END*. One problem with this representation arises from the need to use *abstract* instead of *time indexed* variables. If one needs to reason with both a blend of time-sliced and interval data, then dependencies will exist between the *abstract* variables and the *time-indexed* ones.

The semantics of such arcs and the deployment transformations (conversion to Bayesian network) thereof is not clear from the paper. Presumably, if, in the

MTBN, an *abstract* variable was dependent on a *time indexed* variable then, in the deployed graph, the *abstract* variable would be dependent on each copy of the *time indexed* variable for each time index. If the *time indexed* variable is dependent on the *abstract* variable then the condition is similar in that each copy of the *time indexed* variable is dependent on the *abstract* variable. This results in high degrees of fan-in and fan-out in the deployed graph leading to excessive number of needed probabilities and high complexity.

10 Current Research and Future Work

The probabilistic temporal network can represent very complicated and traditionally difficult domains. Our research has focused on exploring recurrence and periodicity, temporal spacing between cause and effect, and modeling the time-of-reference. These are traditional problems for temporal models. We are currently focusing our efforts on exploring these and other knowledge engineering issues.

In this paper, we introduced a constraint satisfaction formulation for performing belief revision. This formulation needs to be extended to perform belief updating (finding the most likely state of a given interval-RV pair or temporal aggregate). The constraint set needs to be enhanced to take better advantage of the structure imposed by our network structure.

Performing belief revision is in general NP-hard. To address this, we introduced the generalized temporal polytree, which, because of the polytree structure of its dependencies, allows polynomial time belief revision. We are currently investigating practical domains for which the GTP is tenable. The question also remains as to what exactly the maximal tractable class of PTNs is.

Overlapping intervals in a temporal aggregate are troublesome. We allow overlapping intervals so that events happening over intervals can be expressed. For example if a switch could be on from 1000 to 1030 or 1015 to 1045, this could be represented as $\{([1000, 1030], S_0), ([1015, 1045], S_1)\}$ where S_0 and S_1 are random variables for the switches position. S_1 would be conditioned on S_0 to prevent the switch from being on over both intervals. The problem arises in that now the switch could be considered both on and off in the interval [1015, 1030]. Originally, this wasn't considered a problem as the temporal causal relation (TCR) resolved any ambiguity from the perspective of the caused process. One possibility is to make the interval itself random. For example $\{([1000, 1030], S_0), ([1015, 1045], S_1)\}$ might become $\{(I, On)\}$ where $P(I = [1000, 1030]) = P(S_0|\dots)$ and $P(I = [1015, 1045]) = P(S_1|\dots)$. This gets us to only one interval, however now there are two sorts of probabilities to deal with when doing computation.

Our work to date has been within the discrete realm. Future research will focus on modeling continuous domains. Using continuous, rather than discrete, sets of states (Σ) in temporal aggregates is straight-forward. For example, we might have a TA, $Temp = (T_T, \Sigma_T)$ where

$$T_T = \{([0000, 0100], t_1), \dots, ([2300, 2400], t_{24})\}$$

and $\Sigma_T = \mathcal{R}$. *Temp* models changes in the peak temperature over the course of a day. We could have a second TA, *NitDay* = (T_N, Σ_N) where

$$T_N = \{([0000, 0700], n_1), ([0700 - 1900], n_2), ([1900, 2400], n_3)\}$$

and $\Sigma_N = \{\text{night}, \text{day}\}$. With these two TAs, we would like to model peak temperature changing over the course of the day. Temperature during a given hour is dependent on whether or not it is day or night, on the temperature during the previous hour, and on the rate of change between the previous two hours. Constructing the network structure is trivial (see Figure 16).

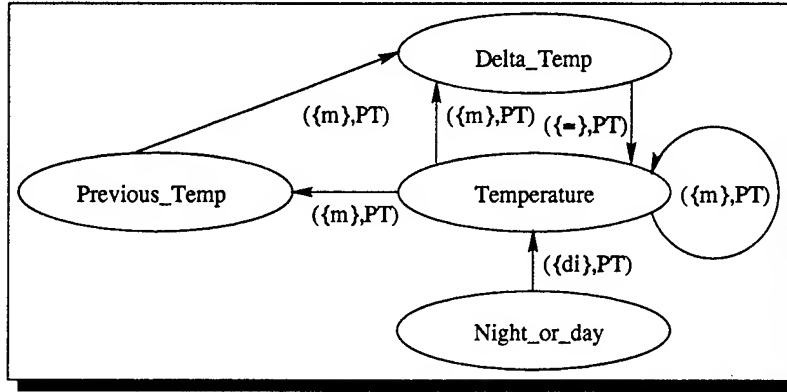


Figure 16: A probabilistic temporal network modeling peak temperature changing over the course of a day.

The difficulty arises in developing appropriate continuous distribution functions for the domain and representing the causal connection between processes (developing appropriate random variable schema) and the conditional dependencies in the caused process. Also, even with continuity in states, without continuity in time, continuous change can not truly be represented. A potential approach is to use a similar structure as discussed for dealing with overlapping interval in which a continuous density function is used to give the probability distribution over the temporal interval space.

11 Conclusion

We have presented a new knowledge representation for merging time and uncertainty. The technique, the probabilistic temporal network, draws from the independence semantics of Bayesian networks and from the interval algebra. Methods for reasoning over the model have been introduced using techniques from operations research and from Bayesian belief revision. A polynomial time subclass is presented, the generalized temporal polytree.

References

- [1] James F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832-843, 1983.
- [2] James F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23(2):123-154, 1984.
- [3] James F. Allen and Johannes A. Koomen. Planning using a temporal world model. In *Proceedings of the Eight International Joint Conference on Artificial Intelligence*, pages 741-747, August 1983.
- [4] Bruce G. Buchanan and Edward H. Shortliffe. *Rule-Based Expert Systems*. Addison Wesley, 1984.
- [5] Eugene Charniak. Bayesian networks without tears. *AI Magazine*, 12(4):50-63, 1991.
- [6] Eugene Charniak and Solomon E. Shimony. Probabilistic semantics for cost based abduction. In *Proceedings of the AAAI Conference*, pages 106-111, 1990.
- [7] Gregory F. Cooper. Probabilistic inference using belief networks is NP-hard. Technical Report KSL-87-27, Medical Computer Science Group, Stanford University, 1987.
- [8] Rina Dechter, Itay Meiri, and Judea Pearl. Temporal constraint networks. *Artificial Intelligence*, 49:61-95, 1991.
- [9] Christian Freska. Temporal reasoning based on semi-intervals. *Artificial Intelligence*, 54:199-227, 1992.
- [10] Avelino J. Gonzalez and Douglas D. Dankel. *The Engineering of Knowledge-Based Systems: Theory and Practice*. Prentice-Hall, Inc., 1993.
- [11] Steve Hanks, David Madigan, and Jonathan Gavrín. Probabilistic temporal reasoning with endogenous change. In *Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, August 1995.
- [12] Jerry R. Hobbs, Mark Stickel, Paul Martin, and Douglas Edwards. Interpretation as abduction. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pages 95-103, 1988.
- [13] John R. Josephson and Susan G. Josephson. *Abductive Inference: Computation, Philosophy, Technology*. Cambridge University Press, 1994.
- [14] Keiji Kanazawa. *Reasoning about Time and Probability*. Thesis, Brown University, May 1992.
- [15] Uffe Kjaerulff. dhugin: A computational system for dynamic time-sliced bayesian networks. *International Journal of Forecasting, Special Issue on Probability Forecasting*, to appear, 1995.
- [16] Gerard Ligozat. Weak representations of interval algebras. In *Proceedings of the AAAI Conference*, pages 715-720, 1990.
- [17] Zohar Manna and Amir Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: specification*. Springer-Verlag, New York, 1992.
- [18] Drew V. McDermott. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101-155, 1982.
- [19] George L. Nemhauser, A. H. G. Rinnooy Kan, and M. J. Todd, editors. *Optimization: Handbooks in Operations Research and Management Science Volume 1*, volume 1. North Holland, 1989.
- [20] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [21] Yun Peng and James A. Reggia. *Abductive Inference Models for Diagnostic Problem-Solving*. Springer-Verlag, 1990.

- [22] Eugene Santos, Jr. On the generation of alternative explanations with implications for belief revision. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 337–347, 1991.
- [23] Eugene Santos, Jr. On modeling time and uncertainty for diagnosis through linear constraint satisfaction. In *Proceedings of International Congress on Computer Systems and Applied Mathematics Workshop on Constraint Processing*, pages 93–106, 1993.
- [24] Eugene Santos, Jr. Unifying time and uncertainty for diagnosis. *Journal of Experimental and Theoretical Artificial Intelligence*, 8:75–94, 1996.
- [25] G. A. Shafer. *Mathematical Theory of Evidence*. Princeton University Press, 1979.
- [26] Yoav Shoham. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press, 1988.
- [27] E. H. Shortliffe. *Computer-Based Medical Consultation: MYCIN*. Elsevier, 1976.
- [28] E. H. Shortliffe and B. G. Buchanan. A model of inexact reasoning in medicine. *Mathematical Biosciences*, 23:351–379, 1975.
- [29] Patrick Suppes. *A probabilistic theory of causality*. Acta philosophica Fennica. Fasc. 24. North-Holland Publishing Co., Amsterdam, 1970.
- [30] Marc Vilain and Henry Kautz. Constraint propagation algorithms for temporal reasoning. In *Proceedings of the AAAI Conference*, pages 377–382, 1986.
- [31] Joel D. Young and Eugene Santos, Jr. Introduction to temporal bayesian networks. In Michael Gasser, editor, *Online Proceedings of the 1996 Midwest Artificial Intelligence and Cognitive Science Conference*. URL: <http://www.cs.indiana.edu/event/maics96/Proceedings/Young/maics-96.html>, April 1996.
- [32] L. A. Zadeh. The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets and Systems*, 11:199–227, 1983.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 13 November 1996		3. REPORT TYPE AND DATES COVERED final
4. TITLE AND SUBTITLE PROBABILISTIC TEMPORAL NETWORKS			5. FUNDING NUMBERS	
6. AUTHOR(S) Eugene Santos Jr., Joel D. Young				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology, WPAFB OH 45433-6583			8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/EN/TR96-06	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) <p>Complex real-world systems consist of collections of interacting processes/events. These processes change over time in response to both internal and external stimuli as well as to the passage of time itself. Many domains such as real-time systems diagnosis, story understanding, and financial forecasting require the capability to model complex systems under a unified framework to deal with both time and uncertainty. Current models for uncertainty and current models for time already provide rich languages to capture uncertainty and temporal information respectively. Unfortunately, these semantics have made it extremely difficult to unify time and uncertainty in a way which cleanly and adequately models the problem domains at hand. This is further compounded by the practical necessity of efficient knowledge engineering under such a unified framework. Existing approaches suffer from significant trade offs between strong semantics for uncertainty and strong semantics for time. In this paper, we define and explore a new model, the Probabilistic Temporal Network, for representing temporal and atemporal information while fully embracing probabilistic semantics. The model allows representation of time constrained causality, of when and if events occur, and of the periodic and recurrent nature of processes. A constraint satisfaction formulation is presented for belief revision as well as a polynomial solvable class.</p>				
14. SUBJECT TERMS temporal reasoning, probabilistic networks, uncertainty, artificial intelligence			15. NUMBER OF PAGES 26 34	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	